

APPLICATION FOR UNITED STATES PATENT

**DISTRIBUTING PACKETS AMONG MULTIPLE TIERS OF NETWORK  
APPLIANCES**

By Inventors:

Mark Albert  
7025 Kit Creek Road  
Research Triangle, NC 27709  
A citizen of the United States

Louis Menditto  
7025 Kit Creek Road  
Research Triangle, NC 27709  
A citizen of the United States

Robert Batz  
7025 Kit Creek Road  
Research Triangle, NC 27709  
A citizen of the United States

Chris O'Rourke  
7025 Kit Creek Road  
Research Triangle, NC 27709  
A citizen of the United States

Richard Gray  
7025 Kit Creek Road  
Research Triangle, NC 27709  
A citizen of the United States

Pranav Tiwari  
7025 Kit Creek Road  
Research Triangle, NC 27709  
A citizen of India

Jacob Mark McGuire  
7025 Kit Creek Road  
Research Triangle, NC 27709  
A citizen of the United States

Tzu-Ming Tsang  
7025 Kit Creek Road  
Research Triangle, NC 27709  
A citizen of the United States

Assignee: Cisco Technology, Inc.

VAN PELT AND YI, LLP  
4906 El Camino Real Suite 205  
Los Altos, CA 94022  
Telephone (650) 903-3500

DISTRIBUTING PACKETS AMONG MULTIPLE TIERS OF  
NETWORK APPLIANCES

CROSS REFERENCE TO RELATED APPLICATIONS

This application claims priority to co-pending U.S. Patent Application No.

- 5 09/346,634, filed July 2, 1999 (Attorney Docket No. CISC514) entitled  
DISPATCHING PACKETS FROM A FORWARDING AGENT USING TAG  
SWITCHING which is incorporated herein by reference for all purposes; and co-pending  
U.S. Patent Application No. 09/347,124, filed July 2, 1999 (Attorney Docket  
No.CISC515) entitled CASCADING MULTIPLE SERVICES ON A FORWARDING  
10 AGENT which is incorporated herein by reference for all purposes; and co-pending U.S.  
Patent Application No. 09/347,111, filed July 2, 1999 (Attorney Docket No. CISC516)  
entitled LOAD BALANCING USING DISTRIBUTED FORWARDING AGENTS  
WITH APPLICATION BASED FEEDBACK FOR DIFFERENT VIRTUAL  
MACHINES which is incorporated herein by reference for all purposes; and co-pending  
15 U.S. Patent Application No. 09/347,428, filed July 2, 1999 (Attorney Docket No.  
CISC517) entitled GATHERING NETWORK STATISTICS IN A DISTRIBUTED  
NETWORK SERVICE ENVIRONMENT which is incorporated herein by reference for  
all purposes; and co-pending U.S. Patent Application No. 09/347,122, filed July 2, 1999  
(Attorney Docket No. CISC518 ) entitled HANDLING PACKET FRAGMENTS IN A  
20 DISTRIBUTED NETWORK SERVICE ENVIRONMENT which is incorporated herein  
by reference for all purposes; and co-pending U.S. Patent Application No. 09/347,108,

filed July 2, 1999 (Attorney Docket No. CISC519) entitled SENDING  
INSTRUCTIONS FROM A SERVICE MANAGER TO FORWARDING AGENTS ON  
A NEED TO KNOW BASIS which is incorporated herein by reference for all purposes;  
and co-pending U.S. Patent Application No. 09/347,126, filed July 2, 1999 (Attorney  
5 Docket No. CISC520) entitled DISTRIBUTION OF NETWORK SERVICES AMONG  
MULTIPLE SERVICE MANAGERS WITHOUT CLIENT INVOLVEMENT, filed July  
2, 1999 which is incorporated herein by reference for all purposes; and co-pending U.S.  
Patent Application No. 09/347,034, filed July 2, 1999 (Attorney Docket No. CISC521)  
entitled INTEGRATING SERVICE MANAGERS INTO A ROUTING  
10 INFRASTRUCTURE USING FORWARDING AGENTS which is incorporated herein  
by reference for all purposes, and co-pending U.S. Patent Application No. 09/347,048,  
filed July 2, 1999 (Attorney Docket No. CISC522) entitled SYNCHRONIZING  
SERVICE INSTRUCTIONS AMONG FORWARDING AGENTS USING A SERVICE  
MANAGER which is incorporated herein by reference for all purposes; and co-pending  
15 U.S. Patent Application No. 09/347,125, filed July 2, 1999 (Attorney Docket No.  
CISC527) entitled BACKUP SERVICE MANAGERS FOR PROVIDING RELIABLE  
NETWORK SERVICES IN A DISTRIBUTED ENVIRONMENT which is incorporated  
herein by reference for all purposes; and co-pending U.S. Patent Application No.  
09/347,123, filed July 2, 1999 (Attorney Docket No. CISC528) entitled STATEFUL  
20 FAILOVER OF SERVICE MANAGERS which is incorporated herein by reference for  
all purposes; and co-pending U.S. Patent Application No. 09/347,109, filed July 2, 1999  
(Attorney Docket No. CISC529) entitled NETWORK ADDRESS TRANSLATION  
USING A FORWARDING AGENT which is incorporated herein by reference for all

purposes; and co-pending U.S. Patent Application No. 09/347,036, filed July 2, 1999  
(Attorney Docket No. CISC530) entitled PROXYING AND UNPROXYING A  
CONNECTION USING A FORWARDING AGENT, which is incorporated herein by  
reference for all purposes; and co-pending U. S. Patent Application No. 09/651,436, filed  
5 August 30, 2000 (Attorney Docket No. CISC536) entitled DISTRIBUTED RULE-  
BASED PACKET REDIRECTION, which is incorporated herein by reference for all  
purposes.

### FIELD OF THE INVENTION

The present invention relates generally to networks. More specifically,  
10 distributing packets among multiple tiers of network appliances is described.

### BACKGROUND OF THE INVENTION

As the IP protocol has continued to be in widespread use, a plethora of network  
service appliances have evolved for the purpose of providing certain network services not  
included in the protocol and therefore not provided by standard IP routers. Such services  
15 include NAT, statistics gathering, load balancing, proxying, intrusion detection, and  
numerous other security services. In general, such service appliances must be inserted in  
a network at a physical location where the appliance will intercept all flows of interest for  
the purpose of making its service available.

Figure 1 is a block diagram illustrating a prior art system for providing a network  
20 service. A group of clients 101, 102, and 103 are connected by a network 110 to a group

of servers 121, 122, 123, and 124. A network service appliance 130 is physically located in the path between the clients and the servers. Network service appliance 130 provides a service by filtering packets, sending packets to specific destinations, or, in some cases, modifying the contents of packets. An example of such modification would be  
5 modifying the packet header by changing the source or destination IP address and the source or destination port number.

Network service appliance 130 provides a network service such as load balancing, caching, or security services. In providing security services, network service appliance 130 may function as a proxy, a firewall, or an intrusion detection device. For purposes of  
10 this specification, a network service appliance that acts as a load balancer will be described in detail. It should be noted that the architecture and methods described are equally applicable to a network service appliance that is functioning as one of the other above described devices.

Network service appliance 130 is physically located between the group of servers  
15 and the clients that they serve. There are several disadvantages to this arrangement. First, it is difficult to add additional network service appliances when the first network service appliance becomes overloaded because the physical connections of the network must be rerouted. Likewise, it is difficult to replace the network service appliance with a back up network service appliance when it fails. Since all packets pass through the  
20 network service appliance on the way to the servers, the failure of the network service appliance may prevent any packets from reaching the servers and any packets from being sent by the servers. Such a single point of failure is undesirable. Furthermore, as

networks and internetworks have become increasingly complex, multiple services may be required for a single network and inserting a large number of network service appliances into a network in places where they can intercept all relevant packet flows may be impractical.

5           The servers may also be referred to as hosts and the group of servers may also be referred to as a cluster of hosts. If the group of servers has a common IP address, that IP address may be referred to as a virtual IP address (VIPA) or a cluster address. Also, it should be noted that the terms client and server are used herein in a general sense to refer to devices that generally request information or services (clients) and devices that  
10           generally provide services or information (servers). In each example given it should be noted that the roles of client and server may be reversed if desired for a particular application.

A system that addresses the scalability issues that are faced by network service appliances (load balancers, firewalls, etc.) is needed. It would be useful to distribute  
15           functions that are traditionally performed by a single network element and so that as much function as possible can be performed by multiple network elements. A method of coordinating work between the distributed functions with a minimum of overhead is needed.

Although network service appliances have facilitated the development of scalable  
20           server architectures, the problem of scaling network service appliances themselves and distributing their functionality across multiple platforms has been largely ignored.

Network service appliances traditionally have been implemented on a single platform that must be physically located at a specific point in the network for its service to be provided.

For example, clustering of servers has been practiced in this manner. Clustering has achieved scalability for servers. Traditional multiprocessor systems have relatively low scalability limits due to contention for shared memory and I/O. Clustered machines, on the other hand, can scale farther in that the workload for any particular user is bound to a particular machine and far less sharing is needed. Clustering has also facilitated non-disruptive growth. When workloads grow beyond the capacity of a single machine, the traditional approach is to replace it with a larger machine or, if possible, add additional processors within the machine. In either case, this requires downtime for the entire machine. With clustering, machines can be added to the cluster without disrupting work that is executing on the other machines. When the new machine comes online, new work can start to migrate to that machine, thus reducing the load on the pre-existing machines.

Clustering has also provided load balancing among servers. Spreading users across multiple independent systems can result in wasted capacity on some systems while others are overloaded. By employing load balancing within a cluster of systems the users are spread to available systems based on the load on each system. Clustering also has been used to enable systems to be continuously available. Individual application instances or machines can fail (or be taken down for maintenance) without shutting down service to end-users. Users on the failed system reconnect and should not be aware that they are using an alternate image. Users on the other systems are completely unaffected

except for the additional load caused by services provided to some portion of the users that were formerly on the failed system.

In order to take full advantage of these features, the network access must likewise be scalable and highly available. Network service appliances (load-balancing appliances  
5 being one such example) must be able to function without introducing their own scaling limitations that would restrict the throughput of the cluster. A new method of providing network services using a distributed architecture is needed to achieve this.

It would also be advantageous if such a distributed architecture could provide network services using multiple tiers of network appliances with each tier providing a  
10 different service.

## SUMMARY OF THE INVENTION

Accordingly, distributing packets among multiple tiers of network appliances is disclosed. In one embodiment, workload is balanced in a serial fashion among two or more tiers of servers. Each tier of server may provide a different service. In the example  
5 described below, requests from a client are first filtered through a tier of firewalls before access is allowed to a tier of web servers. The same firewall is used for both inbound and outbound flows.

It should be appreciated that the present invention can be implemented in numerous ways, including as a process, an apparatus, a system, a device, a method, or a  
10 computer readable medium such as a computer readable storage medium or a computer network wherein program instructions are sent over optical or electronic communication lines. Several inventive embodiments of the present invention are described below.

In one embodiment, a network includes a first tier of forwarding agents connected to a first tier of network devices. A second tier of forwarding agents is connected to a  
15 second tier of network devices. A service manager is configured to receive a packet from one of the forwarding agents; determine the tier of the forwarding agent; and send an instruction to the forwarding agent directing the forwarding agent to forward the packet to a network device connected to the tier of forwarding agents that includes the forwarding agent.

20 In one embodiment, a service manager configured to distribute packets to multiple tiers of forwarding agents includes a network interface configured to receive packets

from a first tier of forwarding agents connected to a first tier of network devices and a second tier of forwarding agents connected to a second tier of network devices. A processor is configured to determine the tier of a sending forwarding agent that sends a packet; and send an instruction to the sending forwarding agent directing the sending forwarding agent to forward the packet to a network device connected to the tier of forwarding agents that includes the sending forwarding agent.

In one embodiment, a method of distributing packets to multiple tiers of forwarding agents includes receiving packets at a service manager from a first tier of forwarding agents connected to a first tier of network devices and a second tier of forwarding agents connected to a second tier of network devices; determining the tier of a sending forwarding agent that sent a packet; and sending an instruction to the sending forwarding agent directing the sending forwarding agent to forward the packet to a network device connected to the tier of forwarding agents that includes the sending forwarding agent.

In one embodiment, a computer program product for distributing packets to multiple tiers of forwarding agents is embodied in a computer readable medium and includes computer instructions for determining a corresponding tier of a sending forwarding agent that sent a packet received at a service manager from a first tier of forwarding agents connected to a first tier of network devices and a second tier of forwarding agents connected to a second tier of network devices; and sending an instruction to the sending forwarding agent directing the sending forwarding agent to

forward the packet to a network device connected to the corresponding tier of forwarding agents.

In one embodiment, a service manager configured to distribute packets to multiple tiers of forwarding agents includes means for receiving packets from a first tier of forwarding agents connected to a first tier of network devices and a second tier of forwarding agents connected to a second tier of network devices; means for determining the tier of a sending forwarding agent that sends a packet; and means for sending an instruction to the sending forwarding agent directing the sending forwarding agent to forward the packet to a network device connected to the tier of forwarding agents that includes the sending forwarding agent.

These and other features and advantages of the present invention will be presented in more detail in the following specification of the invention and the accompanying figures which illustrate by way of example the principles of the invention.

## BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will be readily understood by the following detailed description in conjunction with the accompanying drawings, wherein like reference numerals designate like structural elements, and in which:

5           Figure 1 is a block diagram illustrating a prior art system for providing a network service.

10           Figure 2A is a block diagram of a network architecture that provides network services without requiring a network service appliance to be physically placed at a node through which all incoming and outgoing packets processed by a group of servers must pass.

            Figure 2B is a block diagram illustrating an architecture for a forwarding agent.

            Figure 2C is a block diagram illustrating an architecture for a service manager.

            Figure 3A is a diagram illustrating how a service manager and a forwarding agent cooperate to establish a connection from a client to a selected real machine.

15           Figure 3B is a diagram illustrating how a forwarding agent routes a SYN ACK returned from a host back to a client.

            Figure 3C is a diagram illustrating how a subsequent data packet from client 304 is routed by forwarding agent 302 to host 306.

Figure 4 is a diagram illustrating a network that includes two forwarding agents and two service managers.

Figure 5 is a diagram illustrating how a service manager provides instructions to two separate forwarding agents for handling a connection.

5        Figure 6 is a diagram illustrating a fixed affinity.

Figure 7 is a diagram illustrating a wildcard affinity.

Figure 8A is a diagram illustrating a service message header.

Figure 8B is a diagram illustrating a segment header.

Figure 8C is a diagram illustrating a security message segment.

10       Figure 9A is a diagram illustrating an affinity update wildcard message.

Figure 9B illustrates a fixed affinity update message that is sent by a service manager to a forwarding agent to add a fixed affinity to the receiver's affinity cache or delete a fixed affinity that is stored in the receiver's affinity cache.

Figure 9C is a diagram illustrating an affinity update-deny message.

15       Figure 9D is a diagram illustrating an interest match message for either a wildcard affinity or a fixed affinity.

Figure 9E is a diagram illustrating an IP packet only message.

Figure 10A is a diagram illustrating an affinity identifier segment.

Figure 10B is a diagram illustrating an affinity service precedence segment.

Figure 10C is a diagram illustrating a service manager interest data segment.

Figure 10D is a diagram illustrating a forwarding agent interest data segment.

5        Figure 10E is a diagram illustrating an identity information segment that is used to identify the sender of a service message.

Figure 10F is a diagram illustrating a NAT (Network Address Translation) action segment.

Figure 10G is a diagram illustrating a sequence number adjust action segment.

10       Figure 10H is a diagram illustrating an advertise action segment.

Figure 10I is a diagram illustrating an interest criteria action.

Figure 10J is a diagram illustrating an action list segment.

Figure 11 is a flowchart illustrating a process implemented on a forwarding agent for handling IP packets.

15       Figure 12 is a flowchart illustrating a process implemented on a forwarding agent to determine whether a packet matches a wildcard affinity.

Figure 13 is a flow chart illustrating a process implemented on a forwarding agent for handling a wildcard affinity received from a service manager.

Figure 14 is a flow chart illustrating the process implemented on a forwarding agent for checking fixed affinities or wildcard affinities stored in an affinity data structure  
5 for the purpose of removing expired affinities.

Figure 15 is a flowchart illustrating a process running on a forwarding agent for recording statistics about a packet.

Figure 16A is a table illustrating the information and dispatch flags and the forwarding address field for a fixed affinity.

10 Figure 16B is a table illustrating different values for the NAT address field and the information and dispatch flags for fixed affinities that specify network address translation as an action.

Figure 17 is a block diagram illustrating a multi-tiered load balancing scheme.

Figure 18A is a block diagram illustrating a data structure in the service manager  
15 that is used to keep track of forwarding agents in the first tier.

Figure 18B is a block diagram illustrating a data structure used to keep track of the forwarding agents in a second tier.

Figure 18C is a block diagram illustrating a data structure to keep track of the appliances in the first tier.

Figure 18D is a block diagram illustrating a data structure for keeping track of the appliances in the second tier.

Figure 18E is a block diagram illustrating a connection object data structure for keeping track of servers assigned to various connections.

5        Figure 19 is a flowchart illustrating a process implemented on the service manager when a packet is received from a forwarding agent.

Figure 20 is a flowchart illustrating a process implemented on the service manager for insuring that outgoing packets are assigned to the same firewall as incoming packets belonging to the same connection.

10

## DETAILED DESCRIPTION

A detailed description of a preferred embodiment of the invention is provided below. While the invention is described in conjunction with that preferred embodiment, it should be understood that the invention is not limited to any one embodiment. On the contrary, the scope of the invention is limited only by the appended claims and the invention encompasses numerous alternatives, modifications and equivalents. For the purpose of example, numerous specific details are set forth in the following description in order to provide a thorough understanding of the present invention. The present invention may be practiced according to the claims without some or all of these specific details. For the purpose of clarity, technical material that is known in the technical fields related to the invention has not been described in detail in order not to unnecessarily obscure the present invention.

Figure 2A is a block diagram of a network architecture that provides network services without requiring a network service appliance to be physically placed at a node through which all incoming and outgoing packets processed by a group of servers must pass. Several clients 201, 202, and 203 are connected to a network 210. Network 210 is connected to a group of servers 220 that includes servers 221, 222, and 223. There is no point through which all traffic between devices connected to network 210 and the group of servers 220 must pass. Instead, some traffic from network 210 that is bound for the group of servers passes through a forwarding agent 231 and some traffic between network 210 and group of servers 220 passes through a forwarding agent 232.

In the example shown, forwarding agent 231 is connected to server 221 and server 222 and forwarding agent 232 is connected to server 222 and server 223. Thus, server 222 may communicate with network 210 through either of the forwarding agents, server 221 communicates with network 210 exclusively through forwarding agent 231, and  
5 server 223 communicates with network 210 exclusively through forwarding agent 232. This arrangement may be generalized to include an arbitrary number of servers connected to an arbitrary number of forwarding agents with individual servers connected to arbitrary subsets of the forwarding agents.

A service manager 241 and a second service manager 242 also communicate with  
10 the forwarding agents. The service managers provide the decision making capability that is required to provide a network service such as load balancing. The service managers send specific instructions to each of the forwarding agents detailing how certain flows of packets are to be processed. Such packet processing may include simply routing the packet, gathering statistics about the packet, sending the packet to a service manager,  
15 sending a notification that the packet has been seen to a service manager, modifying the packet, or using a special method such as tunneling or tag switching to send the packet to a destination other than the destination specified by the destination IP address included in the packet header. It should also be noted that forwarding agents in other embodiments also modify other aspects of packets, including packet source and destination addresses  
20 and port numbers and, in some instances, packet data.

The service managers communicate with the forwarding agents to give the agents instructions relating to how to handle packets for various flows that are routed through

the forwarding agents. It is useful at this point to review certain terminology used herein relating to connections and flows.

As used in this specification, a connection consists of a set of flows. A flow is a set of related packets sent between two end stations. A flow may be identified with layer 3 and layer 4 parameters, depending on the protocol being used. For example, for TCP and UDP, a flow is identified by five parameters: the source and destination IP addresses and port numbers and the protocol. For ICMP, flows are defined by three parameters: the source and destination IP addresses and the protocol.

TCP connections will be described in detail in this specification. It should be appreciated that the techniques disclosed apply to other types of connections as well. TCP connections are defined by a 5-tuple that includes the source and destination IP addresses, the source and destination port numbers, and an identification of the protocol that applies to the packet. The source and destination IP addresses and ports for packets going in one direction between the devices are reversed for packets going in the opposite direction. That is, when the direction that a packet is traveling is reversed, the source becomes the destination and the destination becomes the source. Packets flowing in one direction of a connection are in the same flow.

A connection transfers data between applications on two machines having IP addresses and the applications correspond to port numbers. If the protocol is set by convention to be a certain protocol such as TCP, then a protocol identifier may not be required. The 4 remaining numbers, the source and destination IP addresses, and the

source and destination port numbers, are sometimes referred to as a quad. In this specification, the 5-tuple that includes the source and destination IP addresses, the source and destination port numbers and a protocol identification will be referred to as an affinity key. Each unique affinity key thus defines a flow in one direction of a

5 connection. If the source and destination IP addresses and port numbers are reversed for a single affinity key, then it becomes an affinity key that corresponds to a flow in the opposite direction for the same connection. In general, a flow may be identified by a source IP address and destination IP address, by a source IP address, destination IP address and protocol, by a quad, by an affinity key 5-tuple, by only a source and

10 destination IP address or by other information available in a packet header. The term, "flow identifier" is intended to refer to any such method of identifying a flow.

Affinity keys are used by the service managers to identify flows passing through forwarding agents which are to be handled by the forwarding agents in a certain manner. Forwarding agents can accomplish their required tasks with only limited processing

15 capability. Forwarding agents need not determine how to handle certain flows or make decisions such as load balancing or security decisions relating to the flows. The service manager performs those functions and forwards specific instructions to forwarding agents detailing exactly what actions are to be taken for each flow. Instructions for how to handle packets are specified for each flow by the service managers using an affinity key.

20 A specific affinity key that is sent to a forwarding agent together with instructions detailing how packets for flows specified by the affinity key are to be handled is referred to as a fixed affinity.

In addition to specifying instructions for each flow, service managers must also obtain information about each new flow from the forwarding agents. For example, when a service manager provides load balancing through a set of forwarding agents, the service manager uses fixed affinities to provide specific instructions to the forwarding agents

5 detailing where packets for each load balanced flow are to be forwarded. In addition to providing those specific instructions, the service manager also provides general instructions to each forwarding agent that specify which new flows the service manager is interested in seeing. These general instructions are provided using wildcard affinities. Wildcard affinities, which are described in detail below, specify sets of flows that are of

10 interest to a service manager. In one embodiment, this is done by specifying subnet masks that determine sets of source and destination IP addresses that will be forwarded to a service manager. In addition, ports or sets of ports and protocol may be specified in wildcard affinity as well. As is described further below, the use of wildcard affinities enables separate service managers to be configured to provide services for different sets

15 of flows. Each service manager specifies the flows of interest to it and other service managers handle other flows. In this manner, service managers can be configured in parallel to share load.

Thus, service managers use wildcard affinities to specify flows for which they may be providing service and forwarding agents transfer packets for new flows to the

20 appropriate service manager. Once a service manager determines how a certain flow is to be handled, the service manager sends a fixed affinity to each forwarding agent. The fixed affinity overrides the wildcard affinity stored in the forwarding agent that instructs

the forwarding agent to forward packets to the service manager with specific instructions for the specific flow specified by an affinity key in the fixed affinity.

In the case of load balancing, service managers send wildcard affinities to forwarding agents. The wildcard affinities specify destination IP addresses that  
5 correspond to virtual IP addresses of server clusters that are to be load balanced by the service manager. The forwarding agents then forward new packets sent to those virtual IP addresses to the appropriate service manager. The service manager selects a server from the server cluster and then the service manager sends a fixed affinity to each forwarding agent that instructs the forwarding agent to forward packets for that specific  
10 flow to the selected server in the cluster. Forwarding agents may also forward packets for purposes other than load balancing. Packets may be forwarded to real IP addressees as well as virtual IP addresses.

In one embodiment, each forwarding agent is implemented on a router . In other embodiments, forwarding agents may be implemented on switches or other network  
15 devices and may be implemented on a coprocessor in a device that also performs another network function. When implemented on a router, the power of this architecture becomes clear. By infusing each router with a limited functionality provided by the forwarding agent, the service managers are able to provide network services without physically being inserted at the various points in the network where those services must  
20 be provided. The physical presence of each of the routers at those points is sufficient to enable network services to be provided. This contradicts the conventional wisdom regarding the restriction that all traffic inbound for a server cluster must pass through a

single load-balancing engine. The combination of fast forwarding agents (be they 'routers' or IP-aware 'switches') and service managers (to provide synchronization and control) eliminates the scalability limitations of the past.

5 This specification will refer in detail to forwarding agents implemented on routers for the purpose of example. It should be remembered that forwarding agents may also be implemented on other devices and that the same or similar advantages may be realized.

10 The service managers send wildcard affinities to each of the forwarding agents that direct the forwarding agents to process packets that match the wildcard affinities in a certain manner. For example, a service manager may request to be notified when certain packets are received by the routers that include the forwarding agents. When a packet that matches such an instruction is received, the forwarding agent notifies the service manager and the service manager determines what to do with that packet and future packets for the flow based on the network service being provided. Instructions are then sent from the service manager to the forwarding agent at the router that allow the router to process the packets in accordance with the decisions made by the service manager.

20 In addition to specifying that a service manager is to be notified upon receipt of a certain type of packet, wildcard affinities may also specify other actions to be taken. For example, a wildcard may specify an IP address to which packets are to be forwarded without notification to the service manager. Packets may also be copied to a service manager or other device and packets may also be denied or dropped.

It should be noted that the service managers also may be connected to one or more of the servers and may in some cases forward packets received from forwarding agents or received from the network directly to certain servers. However, it is significant that the service managers need not be connected to servers for which they are managing packet traffic. The service manager may accomplish all packet routing through forwarding agents by sending instructions to forwarding agents. It should also be noted that the service managers may also be connected to each other for the purpose of coordinating their instructions or providing backup services.

Figure 2B is a block diagram illustrating an architecture for a forwarding agent. Forwarding agent 250 includes a main processor 252 and a memory 254. Memory 254 may include RAM, ROM, nonvolatile memory such as an EPROM, or a disk drive. Forwarding agent 250 also includes a user interface 256 that allows a user to configure the forwarding agent or monitor the operation of the forwarding agent.

Forwarding agent 250 also includes a service manager interface 258 that allows packets to be sent to and received from a service manager. In addition, the service manager interface allows service managers to send fixed and wildcard affinities to the forwarding agent. In one embodiment, a separate interface is used for the purpose of sending wildcard affinities to forwarding agents using multicast. In other embodiments, a single interface may be provided between the service manger and the forwarding agent.

The forwarding agent also includes a network interface 260 that is used to send and receive packets to and from other devices on the network.

It should be noted that the network interface and the service manager interface may be the same interface in certain embodiments. In such embodiments, all communication between the forwarding agent and the service manager is carried on the same network as packets processed by the forwarding agent.

5 A forwarding agent may be implemented on various network devices. A forwarding agent may be implemented on a network device dedicated to acting as a forwarding agent but the true power of the system is realized when forwarding agents are implemented on network devices that already are included in a network for some other purpose. Forwarding agents may be implemented on routers that already exist at strategic  
10 points in a network for intercepting packets and providing a service using a forwarding agent.

Figure 2C is a block diagram illustrating an architecture for a service manager. Service manager 270 includes a main processor 272 and a memory 274. Memory 274 may include RAM, ROM, nonvolatile memory such as an EEPROM or a disk drive.

15 Service manager 270 also includes a user interface 276 for the purpose of allowing a user to configure the service manager or monitor the operation of the service manager.

Service manager 270 also optionally includes a network interface 278. Network interface 278 allows the service manager to directly forward packets into the network for which it is providing a service. If no network interface is provided, then the service  
20 manager can still forward packets by sending them to a forwarding agent.

A forwarding agent interface 280 is included on the service manager for the purpose of allowing the service manager to send packets and affinities to forwarding agents. Forwarding agent interface 280 may include more than one interface. For example, in one embodiment, a separate interface is used for multicasting wildcard  
5 affinities to all forwarding agents and a separate interface is used for the purpose of unicasting fixed affinities to individual forwarding agents and forwarding packets to individual forwarding agents.

Service manager 270 may also include a service manager interface 282 used to communicate with other service managers. The service manager may communicate with  
10 other service managers for the purpose of providing a fail over scheme of backup service managers. Operational status of service managers may be communicated on the service manager interface and a master service manager may send configuration information about flows being supported through backup service managers so that the backup service managers can function in place of the master service manager should it fail.

15 A service manager may be implemented on a standard microcomputer or minicomputer. In one embodiment a service manager is implemented on a UNIX workstation. A Service manager may also be implemented on other platforms including Windows, an embedded system or as a system on a chip architecture. A service manager also may be implemented on a router.

20 One network service that can be readily provided using the architecture described in Figure 2A is load balancing connections among a set of real machines that are used to

service connections made to a virtual machine. The real machines may also be referred to as hosts and the virtual machine may also be referred to as a cluster of hosts. The following figures describe how a service manager directs forwarding agents to intercept packets for new connections and send them to the service manager. The service manager  
5 then selects a real machine to handle each connection, and directs one or more forwarding agents to forward packets to the selected real machine. Forwarding agents may forward packets using NAT or may use another method of sending packets to the selected real machine.

Figure 3A is a diagram illustrating how a service manager and a forwarding agent  
10 cooperate to establish a connection from a client to a selected real machine. A service manager 300 broadcasts or multicasts a wildcard affinity to all forwarding agents that are listening for wildcard affinities sent by service manager 300. In some embodiments, wildcard affinities may be unicast. A forwarding agent 302 receives the wildcard affinity. In one embodiment, all forwarding agents and service managers register to a  
15 common multicast group so that neither service managers nor forwarding agents need to have any preknowledge of the existence of each other. Thus, a service manager registers its interests with the forwarding agents by multicasting wildcard affinities to the multicast group. Each wildcard affinity provides a filter which recognizes general classes of packets that are of interest.

20 As an example, client 304 may wish to establish a TCP connection with a virtual machine having a virtual IP address. It should be noted that other types of connections may also be established. To establish the TCP connection, client 304 sends a SYN

packet with a destination address corresponding to the virtual IP address. The SYN packet is received by forwarding agent 302. Forwarding agent 302 determines that the destination address of the SYN packet matches the wildcard affinity broadcast by service manager 300. The action included in the broadcast wildcard affinity specifies that all  
5 packets matching the wildcard affinity are to be forwarded to the service manager. Therefore, forwarding agent 302 forwards the SYN packet to service manager 300.

Service manager 300 receives the SYN packet from the forwarding agent. It should be noted that, in one embodiment, forwarding agent 302 encapsulates the SYN packet in a special system packet when the SYN packet is sent to the service manager.  
10 Service manager 300 receives the SYN packet and processes the packet according to whatever service or services are being provided by the service manager. In the example shown, service manager 300 is providing load balancing between a first host 306 and a second host 308. Together, host 306 and host 308 comprise a virtual machine that services the virtual IP address that is the destination of the SYN packet sent by client 304.  
15 Service manager 300 determines the host that is to receive the SYN packet and that is to handle the connection initiated by the SYN packet. This information is included in a fixed affinity. The SYN packet is encapsulated with the fixed affinity and sent back to forwarding agent 302.

The fixed affinity sent to the forwarding agent 302 may include an action that  
20 directs the forwarding agent to dispatch the SYN packet directly to host 306. The action included in the fixed affinity may also direct the forwarding agent to translate the destination address of the packet to the IP address of host 306 and the packet may be

5 routed to host 306 via one or more hops. In addition, as described below, tag switching may also be used to send the packet to the host that is selected by the service manager using its load balancing algorithm.

Thus, the SYN packet is directed to the host selected by service manager 300  
5 without service manager 300 being inserted into the path of the packet between the hosts which comprise virtual machine 310 and client 304. The service manager broadcasts a wildcard affinity to all forwarding agents potentially in that path and the forwarding agents forward SYN packets to the service manager whenever a client establishes a new  
10 connection. The service manager then returns the SYN packet with a fixed affinity that directs the forwarding agent how to forward that SYN packet as well as future packets sent in the flow from the client to the virtual machine. The forwarding agent then sends the SYN packet on to the selected host using network address translation (NAT), tag switching, or some other method.

Figure 3B is a diagram illustrating how a forwarding agent routes a SYN ACK  
15 returned from a host back to a client. A service manager 300 broadcasts a wildcard affinity to a forwarding agent 302. The wildcard affinity matches packets with a source IP address matching either host 306 or host 308 which implement virtual machine 300. When host 306 sends a SYN ACK packet back to client 304, the SYN ACK travels through forwarding agent 302. Because of the wildcard affinity that matches the source  
20 IP address of host 306, forwarding agent 302 encapsulates the SYN ACK packet and sends it to service manager 300. Service manager 300 then identifies the SYN ACK as the SYN ACK corresponding to the SYN that was sent by the client shown in Figure 3A

and sends the SYN ACK together with a fixed affinity to forwarding agent 302. The fixed affinity may include an action that directs the forwarding agent to replace the source IP address of host 306 with the virtual IP address of virtual machine 310 before forwarding the SYN ACK packet on to client 304.

5           Thus, Figures 3A and 3B show how a forwarding agent intercepts a SYN packet from a client and translates the destination IP address from the destination IP address of a virtual machine to the destination IP address of a specific host. The specific host is determined by the service manager using a load balancing algorithm. The forwarding agent does not include logic that performs load balancing to determine the best host. The forwarding agent only needs to check whether the incoming SYN packet matches a fixed  
10           affinity or a wildcard affinity broadcast to the forwarding agent by the service manager.

          The SYN packet is forwarded to the service manager and the service manager returns the SYN packet to the forwarding agent along with a fixed affinity that includes an action which specifies how the forwarding agent is to handle the SYN packet. When a  
15           SYN ACK is returned by the host, the forwarding agent again finds a wildcard affinity match and forwards the SYN ACK packet to the service manager. The service manager returns the SYN ACK packet to the forwarding agent along with a second fixed affinity that instructs the forwarding agent how to handle packets in the flow back from the host the client.

20           The first fixed affinity from the service manager includes an affinity key that corresponds to the flow from the client to the host and the second fixed affinity sent from

the service manager to the forwarding agent contains an affinity key that corresponds to the flow from the host back to the client. Future packets in either flow sent from the client or the host match the affinity key in one of the fixed affinities and are handled by the forwarding agent according to the action contained in the fixed affinity. It is no longer necessary to forward such packets to the service manager. In some applications, the forwarding agent may continue to forward data about the packets to the service manager so that the service manager can monitor connections or maintain statistics about network traffic.

Figure 3C is a diagram illustrating how a subsequent data packet from client 304 is routed by forwarding agent 302 to host 306. Client 304 sends a data packet to forwarding agent 302. Forwarding agent 302 has stored the fixed affinity corresponding to the flow from the client to the host in a fixed affinity database 303. Forwarding agent 302 notes the match of the 5-tuple of the data packet with an affinity key in the fixed affinity database and then forwards the data packet according to the action defined in that fixed affinity. In this example, the action defined is to translate the destination IP address of the client from the virtual IP address of virtual machine 310 to the IP address of host 306. In addition to forwarding the data packet, the affinity found by the forwarding agent also includes an action that requires the forwarding agent to send an affinity packet to service manager 300 that includes data about the packet for the purpose of service manager 300 gathering statistics about network traffic.

The examples shown in Figure 3A through Figure 3C illustrate how the first packet sent in both flows of a new connection are forwarded to the service manager by

the forwarding agent. The service manager then directs the forwarding agent to handle the packets in a certain manner by sending fixed affinities to the forwarding agent for each flow and specifying actions to be performed on the packets. In the example shown, the action involves translating the destination IP address from the client to a specific host  
5 IP address and translating the source IP address in packets from the host to a virtual IP address. Other actions may be defined by fixed affinities including translating other IP addresses, translating port numbers or dispatching packets to other machines. Some of these other actions are described below.

Figure 4 is a diagram illustrating a network that includes two forwarding agents  
10 and two service managers. A first client 402 and a second client 404 send packets through a network or internetwork 406 that eventually reach a subnetwork that includes a first forwarding agent 410, a second forwarding agent 412, a first service manager 420, and a second service manager 422. In the examples shown, the service managers communicate with the forwarding agents and with each other over the same physical  
15 network that is used to send packets. In other embodiments, a separate physical connection may be provided between service managers for the purpose of coordinating service managers and providing back up service managers and a separate connection may be provided between the service managers and the forwarding agents for the purpose of multicasting wildcard affinities or, in some embodiments, for sending fixed affinities and  
20 returning packets to forwarding agents.

In general, the service managers may communicate amongst themselves and with the forwarding agents in any manner appropriate for a particular system. The forwarding

agents each are connected to a first server 430, a second server 432 and other servers up to an nth server 440. These servers may represent one or more virtual machines. Packets from the clients may be routed through either forwarding agent 410 or forwarding agent 412. In fact, packets corresponding to the same connection or flow may be routed at different times through different forwarding agents. To cope with this situation, the service managers multicast wildcard affinities to both forwarding agents. When either forwarding agent first receives a packet for a flow, that forwarding agent forwards the packet to the manager that has requested the packet using a wildcard affinity so that the service manager can provide the forwarding agent with the fixed affinity that defines how to handle the packet.

Figure 5 is a diagram illustrating how a service manager provides instructions to two separate forwarding agents for handling a connection. A client 500 sends a SYN packet to a first forwarding agent 502. Forwarding agent 502 has previously received a wildcard affinity from a service manager 504 on a dedicated connection on which service manager 504 multicasts wildcard affinities to forwarding agents. As a result of the wildcard match, forwarding agent 502 encapsulates the SYN packet and forwards it to service manager 504. Service manager 504 receives the SYN packet and returns it to forwarding agent 502 along with a fixed affinity specifying an action to be performed on the packet. The action defined in this example is translating the destination IP address of the packet from a virtual IP address to the IP address of a host 506. Hosts 506 and 507 together implement a virtual machine 510.

Host 1 receives the SYN packet from forwarding agent 1 and returns a SYN ACK packet back to client 500. However, for some reason, the SYN ACK packet from host 1 is routed not through forwarding agent 502, but instead through forwarding agent 512.

Forwarding agent 512 receives the SYN ACK and notes that it matches a wildcard

5 affinity corresponding to the flow of packets from host 506 to client 500. Forwarding agent 512 encapsulates the SYN ACK packet and sends it to service manager 504.

Service manager 504 defines an action for the SYN ACK packet and includes that action in a second fixed affinity which it sends along with the encapsulated SYN ACK packet

10 back to forwarding agent 512. Forwarding agent 512 then sends the SYN ACK packet on to client 500 where it is processed.

At this point, forwarding agent 502 has a fixed affinity for the flow from client 500 to the hosts and forwarding agent 512 has a fixed affinity for the flow from the hosts back to client 500. Each forwarding agent continues to handle flows without fixed affinities using the wildcard affinities. The service manager acts as a point of

15 synchronization between the forwarding agents when the forwarding agents handle common flows.

Client 500 then sends a data packet which happens to be routed through forwarding agent 512 and not forwarding agent 502. Forwarding agent 502 has received the fixed affinity that provides instructions on how to deal with packets in the flow from

20 client 500 to virtual machine 510. However, forwarding agent 512 has not yet received that fixed affinity. Forwarding agent 512 has received a wildcard affinity previously multicast by the service manager. Therefore, forwarding agent 512 detects a wildcard

affinity match for the data packet and encapsulates the data packet and sends it to service manager 504.

Service manager 504 receives the data packet and notes that the data packet matches the previously defined first fixed affinity which was sent to forwarding agent 502. Service manager therefore does not run the load balancing algorithm again to determine where to route the data packet, but instead returns the first fixed affinity to forwarding agent 512 along with the data packet. Forwarding agent 512 receives the data packet and the fixed affinity and then has the same instructions as forwarding agent 502 for handling that data packet and other packets in the flow from client 500 to virtual machine 510. Forwarding agent 512 therefore translates the destination IP address of the data packet to the IP address of host 506 and forwards the packet on to host 506.

Thus, as long as wildcard affinities are received by each forwarding agent, the service manager is able to provide fixed affinities to each forward agent whenever a fixed affinity is required to provide instructions to handle packets for a given flow. Once a fixed affinity is defined for a flow, the same fixed affinity is provided to any forwarding agent that returns a packet to the service manager as a result of a wildcard match.

To provide a load balancing service for HTTP, a service manager sends a pair of wildcard affinities (one for each direction of flow to and from a virtual machine) to a multicast group that includes each available router in a network. The wildcard affinities specify a protocol and also indicate an exact match on the IP Address and HTTP port number for the virtual machine and an IP address and mask combination that identifies

the client population that is serviced by the service manager. The client population serviced by the service manager is referred to as the client domain of the service manager. If multiple service managers are used, then each service manager may be configured to service a different client domain.

5           For example, if the majority of traffic is coming from a small number of firewalls, whereby the same foreign IP address is shared by many different clients, all those affinities can be assigned by one service manager. Thus, traffic from large sites can be isolated from other traffic and assigned to a different service manager.

10           Thus, the architecture is scalable and service managers may be added to handle client domains as needed. The set of clients serviced by each service manager can be changed by canceling the wildcards that each service manager has broadcast to forwarding agents and sending new wildcards specifying the new client domain.

15           When multiple service managers are included, it is important that the client domains specified by service managers performing the same service do not overlap. The task of assigning affinities for each client domain is centralized by the service manager serving that domain so all packets for a given flow are controlled by a single service manager. For example, if duplicate SYN packets are sent by a client, both should be directed to the same service manager and assigned the same fixed affinity. If the packets were directed to different service managers, then the service manager load balancing  
20           algorithms might assign different real machines to handle the connections as a result of the network being in a different state when the second SYN packet arrived. In addition,

UDP unicasts from the same client must be assigned the same affinity and related connections (e.g., FTP control and data connections) must be assigned the same affinity.

Once the forwarding agents have received fixed affinities, packets intercepted that match a fixed affinity are processed as instructed in the set of actions specified in the fixed affinity. If a matching fixed affinity is not found, the packet is compared against the wildcard affinities to find manager(s) that are interested in this type of packet. If no appropriate Wildcard Affinity is found, normal IP routing occurs. Generally, a manager uses the wildcard affinity to be informed of flows it may be interested in. Once a manager has determined how a flow should be handled, it usually sends a fixed affinity so that the processing of subsequent packets for that flow can be offloaded to the forwarding agent. In some cases actions for certain flows can be predetermined by the service manager without seeing packets from the flow. In such cases, the actions may be specified in a wildcard affinity and no message need be sent to the service manager and no fixed affinity need be generated. The service manager may specify that it is still to receive certain packet types after a fixed affinity is sent by including an optional action interest criteria message segment with the fixed affinity.

In the load-balancing case, a fixed affinity is used to identify the server that is to receive this particular flow whereas a wildcard affinity is used to define the general class of packets for which load balancing is to be performed (all those matching the cluster address and port number for the clustered service) and to identify the manager that is to make the balancing decision for flows that match the wildcard affinity.

## Fixed Affinities

Figure 6 is a diagram illustrating a fixed affinity 600. Fixed affinity 600 matches only one flow through a network. As described above, a flow is defined by an affinity key, which is a unique 5-tuple that spans the packet headers:

5        **IP Header:**

Protocol Type (e.g., UDP or TCP)

Source IP Address

Destination IP Address

**TCP or UDP Header:**

10       Source Port

Destination Port

It should be noted that if the protocol being used is not TCP or UDP, then the ports in the affinity key may be set to 0.

Fixed affinity 600 includes an affinity key 602. In addition, fixed affinity 600  
15 contains information that dictates how a forwarding agent is to process packets that match the affinity key, and how the forwarding agent is to manage the affinity.

A dispatch flag 604 indicates whether the packet is to be dispatched to the forward IP address included in the fixed affinity. Setting the dispatch flag indicates that the packet is to be forwarded to a forward IP address 608 that is provided in the fixed  
20 affinity. The difference between dispatched and directed traffic is that dispatch traffic is forwarded directly from a forwarding agent to a specific server without translating the

destination IP address of the packet. In other words, if a packet is dispatched, then the packet destination address is not used to forward the packet. Instead, a forwarding address contained in an affinity is used to forward the packet. If the connection is not dispatched but directed by the forwarding agent, then the packet IP destination must be  
5 translated using NAT if the packet is redirected to a specific server.

If forward IP address 608 is zero, then the packet is dropped after processing statistics as indicated by an information flag 606. Not setting the dispatch flag indicates that the packet is to be forwarded based on the address provided in the packet IP header.

Information flag 606 indicates whether or not statistics are to be gathered for  
10 packets forwarded using the fixed affinity. If the Information flag is set, statistics are updated for the forward IP address. In one embodiment, the statistics kept include:

1. total bytes for all packets matching the forward IP address
2. total packets matching the forward IP address

Statistics for packets and bytes matching the affinity may be kept regardless of the  
15 setting of the Information flag.

Fixed affinity 600 also includes a time to live 610. Time to live 610 specifies the number of seconds before the fixed affinity should be timed-out from a fixed affinity cache maintained by a forwarding agent. If a time to live of 0 is specified, then that means that the fixed affinity is not to be cached by a forwarding agent and if a copy of  
20 the fixed affinity is already in the cache, it should be removed. Thus, service managers

may remove fixed affinities that they have sent to forwarding agents by simply sending copies of those fixed affinities to the forwarding agents with time to live set to 0.

Each fixed affinity sent by a service manager is correlated to a wildcard affinity previously sent by the service manager. If a forwarding agent receives a fixed affinity for  
5 which no supporting wildcard affinity is found, the forwarding agent ignores the fixed affinity and discards it.

### **Wildcard Affinities**

Figure 7 is a diagram illustrating a wildcard affinity 700. Wildcard affinity 700 is a more general form of Affinity that is used by a service manager to register filters with  
10 the forwarding agent(s) that define the range of flows that are of interest to the service manager. Like a fixed affinity, wildcard affinity 700 also includes a dispatch flag 702 and an information flag 704. Wildcard affinity 700 also includes the elements of an affinity key (protocol 706, source IP address 708, destination IP address 712, source port 716, and destination port 718) plus source netmask 710 and destination netmask 714.

15 The netmasks and the source and destination IP addresses are used to specify ranges of addresses covered by the wildcard affinity. The source netmask is ANDed with the source IP address in the wildcard affinity. The source netmask is also ANDed with the source IP address from the packet. If the results of the two operations are equal, then the source IP address of the packet is considered to be in range of the wildcard affinity.  
20 Likewise, the destination netmask is ANDed with the destination IP address in the wildcard affinity. The destination netmask is also ANDed with the destination IP address

from the packet. If the results of the two operations are equal, then the destination IP address of the packet is considered to be in range of the wildcard affinity. If both the source and the destination IP addresses of the packet are in the range of the wildcard affinity, and the ports and protocols also match, then the packet is said to match the  
5 wildcard affinity. It should also be noted that, in one embodiment, a zero specified for a port or a protocol matches all ports or protocols.

It should be noted that in other embodiments, other methods of specifying ranges for the wildcard affinity are used. For example, in one alternative arrangement, ranges of IP addresses are specified by specifying lower bound and upper bound IP addressees. All  
10 addresses between the two bounds fall within the range of the wildcard affinity. In some applications, multiple ranges may be specified. The method described above is particularly useful for specifying a single address, specifying all addresses in a subnet, or specifying every even or odd address, every fourth address, every eighth address, etc.

For example, to specify a single host of 1.1.1.1, the wildcard affinity include an  
15 IP address of 1.1.1.1 with a netmask of 255.255.255.255. To specify the range of hosts from 1.1.1.0 to 1.1.1.255, the wildcard affinity would include an IP address of 1.1.1.0 with a netmask of 255.255.255.0, indicating that the first three bytes of the IP address must match exactly and that the last byte is to be ignored.

Wildcard affinity 700 also includes a time to live 722. Time to live 772 is used in  
20 the same manner as the time to live for the fixed affinity. Wildcard affinities are deleted by forwarding agents based on the time to live set for the wildcard affinity by the service

manager. The timing of such a deletion need not be exact. In one embodiment, the timing need only be accurate to within two seconds. This same tolerance is for fixed affinities as well. Service managers must refresh each wildcard affinity before its time to live expires in order to continue to receive packets that match the wildcard affinity from forwarding agents. As with the fixed affinity, a wildcard affinity may be deleted by sending a duplicate wildcard affinity with a time to live of 0.

### **Actions**

Thus, fixed affinities specify individual flows and packets and wildcard affinities specify sets of flows to be processed in a special way. Such processing is defined by associating actions with the affinities. Actions defined for the affinities specify the service to be performed by the forwarding agent on behalf of the Manager. For fixed affinities, services specified may include:

- Interest Criteria - a list of packet types that cause a notification to be sent to the service manager.
- Sequence Number Adjustment - a set of deltas and initial sequence numbers by which the TCP sequence numbers and ACK numbers are to be adjusted.
- NAT - provides details for how Network Address Translation is to be performed.

For Wildcard Affinities, applicable actions are:

- Interest Criteria - a list of packet types that cause a notification to be sent to the service manager.
- Advertise - indicates that the destination IP Address in the Wildcard Affinity is to be advertised by the forwarding agent. This may be done by including the destination IP address in routing protocol updates.
- Sequence Number Adjustment - a set of deltas and initial sequence numbers by which the TCP sequence numbers and ACK numbers are to be adjusted.

- NAT - provides details for how Network Address Translation is to be performed.

5 Forwarding agents may not support all possible actions. For example, some forwarding agents may not support NAT. The set of actions that the service manager expects a forwarding agent to support are identified in an action list which may be included with the wildcard affinity. If the forwarding agent does not support one or more of the actions identified in the list, it discards the wildcard affinity and sends a message to  
10 the service manager indicating that it does not support all of the actions in the list. This message is referred to as an affinity update deny message. The service manager then may attempt to send a new wildcard affinity that excludes any unsupported actions identified in the affinity update deny message.

### Service Messages

15 Wildcard affinities, fixed affinities, actions, packets, and other messages are sent between service managers and forwarding agents encapsulated in service messages. In one embodiment, messages sent between service managers and forwarding agents are sent using the specific service message format described below. Service messages are sent between service managers and forwarding agents using UDP. Wildcard affinities,  
20 which are sent by service managers, can be multicast to a multicast IP Address and UDP Port known to the service manager(s) and forwarding agent(s), or can be unicast to a particular forwarding agent or service manager. Figure 8A is a diagram illustrating a service message header used in one embodiment. Service message header 800 includes a protocol version 802 and a message type 804. The protocol version identifies the version

of the service protocol supported by the sender. The message type identifies the overall purpose of this message, the base format for the message, and implies the set of optional message segments that may be included in the message.

The following service message types are used:

Message Type
affinity update-wildcard affinity
affinity update-fixed affinity
affinity update-deny
interest match-wildcard affinity
interest match-fixed affinity
IP packet only

5

The affinity update-wildcard affinity message is used to send wildcard affinities from a service manager to forwarding agents. The affinity update-fixed affinity message is used to send fixed affinities. The affinity update-deny message is used to report that an affinity update message has been rejected because required actions included in the affinity update are not supported by the receiver. The interest match-wildcard affinity message is used to report a wildcard affinity match to a service manager and the interest match-fixed affinity message is used to report a fixed affinity match to a service manager. The IP packet only message is used to forward an IP packet.

10

After the service message header, a service message includes one or more message segments. Each message segment begins with its own segment header. Figure

15

8B is a diagram illustrating a segment header. Segment header 810 includes a Required flag 812. Required flag 812 defines whether the sender will allow the rest of the message to be processed even if the segment cannot be processed (either because the receiver does not support the function described by the segment or because the receiver does not understand the segment). The required flag either indicates that the segment may be ignored or that the segment is required. If a required segment cannot be processed, then the entire message that includes the segment is dropped and an error message is returned to the sender. Each segment header is followed by data that is specific to the message segment.

10           The following message segments are used:

Segment Name
Wildcard Affinity
Fixed affinity
Affinity Interest
Service Precedence
Security
Service Manager Interest Data
forwarding agent Interest Data
Identity Info
Action-NAT
Action-Advertise
Action-Sequence Number Adjust
Action-Interest Criteria
Action List
IP Packet

The fixed affinity, wildcard affinity and security segments are described immediately below. The remaining segments are described in detail following a

5 description of the message types that include the segments.

## Security

If security is expected by the receiver, a security message segment immediately follows the service message header. The security message segment contains the expected security sequence. If the receiver does not expect security, the security message segment is ignored (if present) and the message is accepted. Security is generally not required for IP packet only messages. If authentication is successful, the signals are accepted. If the authentication fails, the signal is ignored. Various authentication schemes such as MD5 may be supported. The type of authentication to be used is configured at the senders and receivers, along with a password. If the receiver does not expect authenticated messages, then the security segment may be ignored if it is present and the signal may be accepted whether or not it contains a security segment.

Figure 8C is a diagram illustrating a security message segment. Security message segment 820 includes a security type field and a security data field 824. Security type field 822 describes the type of encoding used for security (i.e., MD5, etc.). Security data field 824 contains the data needed to implement the algorithm identified by the security type field 822.

### **Detailed Message Descriptions**

#### **Wildcard Affinity Update**

Figure 9A is a diagram illustrating an affinity update wildcard message. Affinity update wildcard message 900 is sent by a service manager to a forwarding agent to register or unregister for classes of flows that match the specified sets of flows. It includes a service message header 902 followed by a sequence of message segments. A

security segment 903 is optional, as dictated by the needs of the receiver. A wildcard affinity segment 904 is required, since the purpose of the affinity update wildcard message is to send a wildcard. An action list segment 906 is optional. Its purpose is list the actions that a forwarding agent must support in order to receive the affinity. If the forwarding agent determines that any of the actions are not supported, then it may send an affinity update deny message to the service manager.

An affinity service precedence field 908 is optionally used to specify the precedence of the service being provided. This allows multiple service managers or a single service manager to send wildcard affinities for different services. An affinity backup precedence field 909 is also optionally used to specify the backup precedence of the service manager that sent the affinity. This allows a backup service manager to send wildcard affinities that are ignored until a higher backup service precedence wildcard affinity that corresponds to a primary service manager is deleted. An identity information segment 910 is optionally used to identify the manager. This information may be used, for example, in an error message on the console of the forwarding agent to indicate which service manager had a problem. A service manager interest data segment is optionally used to include data that should be returned to the service manager when an interest match-wildcard affinity message is sent to the service manager as a result of a forwarding agent determining a wildcard affinity match. Finally, one or more action segments are optionally included. The action segments specify actions that are performed on the packets for the purpose of providing a network service. It should be noted that in some embodiments, fields which are described above as optional may become required and

required fields may be optional. This is also generally true of the other message descriptions contained herein.

### **Fixed Affinity Update**

Figure 9B illustrates a fixed affinity update message that is sent by a service manager to a forwarding agent to add a fixed affinity to the receiver's affinity cache or delete a fixed affinity that is stored in the receiver's affinity cache. If the time to live in the fixed affinity segment is non-zero, the affinity is added to the cache (or refreshed, if it already resides there) for the number of seconds specified in the time to live. If time to live is zero, the fixed affinity is removed from the cache if it is found there.

Fixed affinity update message 920 includes a service message header 922. An optional security segment 924 is included as dictated by the needs of the receiver. A fixed affinity segment 926 includes the fixed affinity being sent. An affinity service precedence 928 optionally specifies a service precedence. An affinity backup precedence field 929 is also optionally used to specify the backup precedence of the service manager that sent the affinity. This allows a backup service manager to send affinities that are ignored until a higher backup service precedence affinity that corresponds to a primary service manager is deleted. One or more action segments 930 are optionally included to specify actions to be performed by the receiver for matching packets. An identity information segment 932 is optionally used to identify the service manager that sent the fixed affinity. A service manager interest data segment 934 is optionally used to include data that should be returned to the service manager when an interest match-wildcard

affinity message is sent to the service manager as a result of a forwarding agent determining a wildcard affinity match. A forwarding agent interest data segment 936 is optionally used to include data that a forwarding agent requested to be returned to it along with a fixed affinity. Finally, an IP packet segment 938 includes an IP packet.

5           Usually, the IP packet segment is an IP packet that was sent to a service manager as a result of a wildcard affinity match and that is being sent back to a forwarding agent along with actions to be performed for the packet. In many implementations, the forwarding agent does not devote resources to storing packets that have matched a wildcard affinity and have been forwarded to a service manager. Therefore, the  
10       forwarding agent sends the packet to the service manager along with an interest match message and the service manager sends the packet back to the forwarding agent with a fixed affinity update. Thus, the service manager stores the packet for the forwarding agent and returns it to the forwarding agent when the forwarding agent needs to execute an action on the packet. This eliminates the need for storage and garbage collection at  
15       the forwarding agent for packets that matched a wildcard affinity and are awaiting instructions from a service manager for handling. In some implementations, the forwarding agents may temporarily store packets that have matched a wildcard affinity. However, it has been found that sending packets to the service manager and having the service manager return packets with fixed affinities simplifies and improves the  
20       performance of the forwarding agent.

#### **Affinity Update-deny**

Figure 9C is a diagram illustrating an affinity update-deny message. An affinity update-deny message is sent by the forwarding agent to a service manager when the forwarding agent receives an affinity update with a required segment that it cannot process (one where the 'Required' flag is set either within the segment header or within the list of segment types from the action list, if one was included). The segments that cannot be processed properly are identified in the action list that is returned with the affinity update-deny message.

Affinity update-deny message 940 includes a service message header 941. An optional security segment 942 is included as dictated by the needs of the receiver. An action list segment 944 includes actions that are not supported by the forwarding agent and that caused the forwarding agent to send the affinity update-deny message. An affinity segment 946 from the original affinity update that prompted the affinity update-deny message is optionally included. An identity information segment 948 is from the original affinity update that prompted the affinity update-deny message is also optionally included. A service manager interest data segment 950 is optionally used to include data that the service manager sent to the forwarding agent for the forwarding agent to send back to the service manager when an interest match-wildcard affinity message is sent to the service manager. The service manager interest data is used by the service manager to help process the message. A forwarding agent interest data segment 952 is optionally used to include data that the forwarding agent requests to be returned to it along with a fixed affinity.

#### **Interest Match (Wildcard affinity or Fixed affinity)**

Figure 9D is a diagram illustrating an interest match message for either a wildcard affinity or a fixed affinity. Interest match message 960 is sent by the forwarding agent to a service manager when an IP packet matches the interest criteria that was sent the last time the matching affinity was refreshed or added in the cache. Interest match message 5 960 includes a service message header 962. An optional security segment 964 is included as dictated by the needs of the receiver. An affinity identifier segment 966 includes the affinity key of the affinity that caused the match, the dispatch and information flags of that affinity, and an interest match field that provides reasons from the interest criteria that caused the match. In one embodiment, a bit vector is used to 10 provide the reasons.

An identity information segment 968 is optionally included from the original affinity update that prompted the interest match message to be sent. A service manager interest data segment 970 is optionally used to include data that the service manager requested when an interest match message is sent to the service manager. A forwarding 15 agent interest data segment 972 is optionally used to include data that a forwarding agent requested to be returned to it along with a fixed affinity. Finally, an IP packet segment is optionally included so that the forwarding agent can send the IP packet that caused the affinity match to the service manager. The IP packet is sent if the corresponding data flag in the interest criteria indicated that the IP Packet should be sent. The IP packet may 20 be sent as a segment of the interest match message or may be forwarded independently in a subsequent IP Packet message, depending on the capabilities of the forwarding agent.

### **IP Packet Only**

Figure 9E is a diagram illustrating an IP packet only message. IP packet only message 980 is sent by a forwarding agent to a service manager or vice versa whenever an IP network packet is sent from one to the other. This can occur in a number of situations, e.g.,:

5 (1) When a forwarding agent needs to send a service manager a packet that could not be included with an interest match message.

(2) When a forwarding agent needs to send a service manager a packet that matched a service manager wildcard affinity.

10 (3) When a service manager needs to send a forwarding agent a packet that it has processed and that needs to be forwarded to the next appliance (or, if there are no other appliances, to its correct destination). Encapsulating IP packets in the IP packet only message avoids loops in the system by signaling the forwarding agent that the packet has already been to the manager and need not be sent there again.

15 IP packet only message 980 includes a service message header 982. An IP Packet segment 984 includes the IP packet. Preferably IP packet only message 980 does not include a security segment, since the flow is essentially just another IP hop and faster forwarding can be achieved without a security segment.

20 The messages sent between forwarding agents and service managers have now been described in some detail. The wildcard affinity segment, the fixed affinity segment, and the security segment have also been described. The remaining message segments are

described in greater detail below in connection with Figures 10A through 10I. It should be noted that each segment includes, in addition to the fields that are shown, a segment header.

Figure 10A is a diagram illustrating an affinity identifier segment. Affinity identifier segment 1000 includes a dispatch flag 1002, an information flag 1004, and an affinity key 1006. These fields are defined the same as they are defined for fixed affinities and wildcard affinities. Affinity identifier segment 1000 also includes an interest mask 1008 that provides reasons from the interest criteria sent by the service manager that caused the match. This gives the service manager notice of what affinity caused the match and also what interest criteria in that affinity caused the match. The interest criteria action specified in an affinity sent by a service manager is described further below.

Figure 10B is a diagram illustrating an affinity service precedence segment. Affinity service precedence segment 1010 includes a search order flag 1012 that specifies the search order for the precedence, i.e., whether a higher priority precedence is represented by a higher or a lower priority number. A precedence value field 1014 actually provides the precedence value. The service precedence enables one or more service managers to provide different services that are executed in sequential order based on the precedence values provided. In this manner, multiple affinities may be specified that match a flow, with each affinity corresponding to a different service that specifies different actions to be performed for packets in the flow. A packet for such a flow may be forwarded to several service managers before it is eventually sent to the client or the

specific server. It should be noted that only the last service manager can dispatch the packet since the packet must be returned by higher priority service managers to the forwarding agent for further processing by lower priority service managers.

Thus, the affinity service precedence allows multiple service managers of  
5 different types to control the same flow. The value of the precedence dictates the order in which the forwarding agent should process affinities if multiple matches occur. When a matching affinity contains an action that requires the packet to be sent to a service manager, the action is honored. When the packet is returned, the forwarding agent processes the affinity contained in the response and continues with the matching affinity  
10 of the next highest precedence.

Figure 10C is a diagram illustrating a service manager interest data segment. Service manager interest data segment 1020 includes an interest data field 1021 that can contain anything that the service manager arbitrarily determines. This is simply data that can be sent by the service manager to the forwarding agent. The forwarding agent returns  
15 the data to the manager with an interest match message when an interest match is determined. Typically, this data is used to index the affinity.

Figure 10D is a diagram illustrating a forwarding agent interest data segment. Forwarding agent interest data segment 1022 includes an interest data field 1023 that can contain anything that the forwarding agent arbitrarily determines. This is simply data that  
20 can be sent by the forwarding agent to the service manager when an interest match is sent to the service manager. The service manager returns the data to the forwarding agent

with any fixed affinity update message that is sent as a result of the interest match.

Typically, this data is used to index the affinity.

Figure 10E is a diagram illustrating an identity information segment that is used to identify the sender of a service message. The identity information may be used for logging and debugging. Identity information segment 1024 includes an IP address field 1025 that contains the IP address of the message sender. A character field 1026 contains the name of the host.

Figure 10F is a diagram illustrating a NAT (Network Address Translation) action segment. NAT action segment 1030 includes fields that specify a source IP address 1032, a source port 1034, a destination IP address 1036, and a destination port 1038 that are to replace the corresponding fields in the packet. The NAT action segment thus specifies that NAT is to be performed on any packet that matches the associated affinity. A NAT action segment can be included with any Wildcard or Fixed affinity sent by a service manager to a forwarding agent. The action is not performed on packets that are forwarded to the service manager. If the packet is forwarded to the service manager, then the packet is not immediately altered. If the service manager sends the packet back to the forwarding agent for forwarding, the action is performed by the forwarding agent at that time, therefore removing the need for the manager to implement that function directly.

Figure 10G is a diagram illustrating a sequence number adjust action segment. Sequence number adjust action segment 1040 specifies that a forwarding agent should adjust sequence numbers and ACK numbers in the TCP packets that match the associated

affinity. A sequence number adjust action segment can be included with any wildcard affinity or fixed affinity sent by a service manager. The sequence number adjust is not performed on packets that are forwarded to the service manager. The action may be performed when the service manager returns the packet back to the forwarding agent for forwarding.

A sequence delta field 1042 specifies the amount by which the sequence number in packets is to be adjusted. An initial sequence number 1044 specifies the lowest sequence number to which the delta is to be applied. An ACK delta field 1046 specifies the amount by which to adjust the ACK number. An Initial ACK number field 1048 specifies the lowest ACK number to which ACK Delta is to be applied. Thus, sequence numbers and ACK numbers in packets can be modified by forwarding agents according to a scheme determined by a service manager. The scheme is sent to the forwarding agents using the sequence number adjust action segment.

Figure 10H is a diagram illustrating an advertise action segment. An advertise action segment is sent by a service manager to a forwarding agent to specify that the destination IP address in an enclosed wildcard affinity is to be advertised by the forwarding agent. That means that the address is included in routing protocol updates, just as if the destination IP address belonged to a device connected to the router. The address advertisement is deleted when the associated wildcard affinity is deleted. By directing a forwarding agent to advertise an address, the service manager can simulate the presence of an network service appliance at the location of the forwarding agent. For example, if the service manager is providing load balancing among a group of hosts, the

service manager would direct a forwarding agent to advertise the virtual IP address of the cluster of hosts. Thus, the virtual IP address can be advertised as if a load balancer at the location of the forwarding agent were advertising the virtual IP address. If a forwarding agent receives a packet destined for the advertised address, but that packet does not  
5 match an affinity (either Full or Wildcard), the packet is dropped. This avoids establishing connections to the forwarding agent for ports that no service manager is supporting.

Advertise action segment 1050 includes an advertise address 1052, which is the address to be advertised by the forwarding agent. A subnet mask 1054 may also be used  
10 for such advertising. If a subnet mask is used, then the IP address and mask combination indicates a subnet to be advertised. The advertise segment can also be used without specifying a subnet mask.

Figure 10I is a diagram illustrating an interest criteria action. Interest criteria action 1060 is sent by a service manager to a forwarding agent to specify that the service  
15 manager is to be informed when certain types of special packets are detected by the forwarding agent. Interest criteria action 1060 includes an interest IP address 1062 and an interest port 1064. The interest IP address and port specify an IP address and port to which the interest match message is to be sent. An interest mask 1066 is bit vector that specifies the types of packets for which the service manager is requesting notification.  
20 The type of packet specified by the bits may be a function of the protocol type specified in the affinity encapsulated with the interest criteria action. For example if the protocol is TCP, then in one embodiment, the bits are interpreted as follows:

Bit 0 = 1 :: FIN

Bit 1 = 1 :: SYN

Bit 2 = 1 :: RST

Bit 3 = 1 :: PSH

5 Bit 4 = 1 :: ACK

Bit 5 = 1 :: URG

Bit 6 = 1 :: Data Present

Bit 7 = 1 :: First Data present

Bit 8 = 1 :: Fragmented packet, and the source/destination IP addresses match

10 Bit 15 = 1 :: All Packets

If the protocol is UDP, then the bits are interpreted as follows:

Bit 6 = 1 :: Data Present

Bit 7 = 1 :: First Data present

Bit 8 = 1 :: Fragmented packet, and the source/destination IP addresses match

15 Bit 15 = 1 :: All Packets

For other protocols, Bit 15 may be set to indicate all packets.

A data flag 1067 uses the same bit code as the interest mask. Whereas the interest mask determines whether the service manager should be forwarded an interest match message, data flag 1067 specifies whether the service manager is to receive a copy of the packet that caused the interest match with the interest match message. If a bit is set, then the forwarding agent is to send the packet as well as the interest match to interest IP address 1062 and interest port 1064. It should be noted that in some embodiments, the forwarding agents may send messages and forward packets to service managers over a

different network so that the interest IP address and interest port may not be used or some other method may be used for specifying where interest match messages and packets should be sent to the service manager.

A copy flag 1068 also uses the same bit code as the interest mask. Each bit  
5 specifies whether a copy of the matching packet is to be forwarded to the server. If the bit is set for the packet type, the forwarding agent sends a copy of the matching packet and refers to a hold flag 1069 to determine what to do with the original packet. Hold flag 1069 also uses the same bit code as the interest mask. Hold flag 1069 determines whether the forwarding agent forwards the packet to the service manager or, if possible,  
10 holds the packet and waits for the service manager to send a fixed affinity that specifies how the packet should be forwarded by the forwarding agent. If the bit is not set for the packet type, then the forwarding agent forwards the packet. If the bit is set, then the forwarding agent holds the packet, if possible. If the packet cannot be held by the forwarding agent for some reason (e.g., lack of storage) then the forwarding agent  
15 forwards the packet to the Manager.

Figure 10J is a diagram illustrating an action list segment. Action list segment 1070 is sent by a service manager to a forwarding agent with wildcard affinities to specify all the actions that must be supported in order for the forwarding agent accept the wildcard affinity. Action list segment 1070 does not specify that the actions are to be  
20 performed. Its purpose is to warn the forwarding agent of the service requirements. The forwarding agent responds with an affinity update-deny and discards a wildcard affinity if the forwarding agent cannot support all the actions in an action list that is provided with the wildcard affinity. Action list segment 1070 includes a first action type 1072. Action

list segment 1070 may also include a second action type 1074 and other action types up to an nth action type 1080.

A service message protocol for sending messages and packets between service managers and forwarding agents has been defined in Figures 6-10J. Each service message includes a service message header that identifies the message type. After the service message header, each service message includes one or more segments, depending on the message type. Each segment begins with a segment header. Using the message types described, service managers can send forwarding agents instructions detailing certain sets of packets that the service manager wants to either to be forwarded to the service manager or to cause an interest match message to be sent to the service manager. Messages are also used to specify actions for certain packets in certain flows.

For example, if a service manager is providing load balancing, the service manager first sends a wildcard affinity update message to a forwarding agent specifying a set of clients that the service manager will load balance. The wildcard affinity may also include an action that directs the forwarding agent to advertise a virtual IP address for a virtual machine that includes all of the load balanced servers. When the forwarding agent intercepts a packet that matches the wildcard affinity, then the forwarding agent sends an interest match message to the service manager. The service manager then determines a server to assign the connection (or the server that has already been assigned the connection) and sends a fixed affinity to the forwarding agent that directs the forwarding agent to dispatch the packet to that server or to use NAT to substitute the server's address in the packet. The service manager also may include an interest criteria in a fixed affinity

that specifies that future packets for the flow should not be sent to the service manager, but that the service manager should be notified if certain types of packets such as a FIN or a FIN ACK are received. At any point, the service manager may cancel a fixed affinity or a wildcard affinity sent to a forwarding agent by sending a fixed affinity or a  
5 wildcard affinity with a time to live of 0.

Thus service managers are able to control affinities and monitor flows using the above defined messages. When a forwarding agent receives a packet, affinities received from service managers are searched first for the one with the highest service precedence. Once a match is determined, the search order defined for that precedence is used to find  
10 another identical Affinity with a better service precedence. If multiple affinities exist with the same best service precedence, they are searched for the one with the lowest backup precedence value.

Service managers manage the storage of affinities on forwarding agents using the time to live portion of the affinity segments. The forwarding agents remove affinities at  
15 intervals specified by the service manager if they have not already been removed at the request of a manager (via an affinity update message with a time-to-live of zero). No affinity is kept for an interval longer than the interval specified by the time-to-live set by the manager (within a tolerance of +/- 2 seconds in one embodiment) so that the manager can reliably assume that the affinities have been cleared at some small time beyond that  
20 interval that accounts for any propagation or processing delays. This simplifies the managing of affinities by the service manager across multiple routers. In some cases, a

forwarding agent may need to ask for an affinity again if more traffic arrives for that affinity after it has been deleted.

The service manager itself stores affinities long enough to allow forwarding agents sufficient time to delete their own copies. If an affinity is allowed to expire at a service manager, it must be kept by the service manager long enough so that the forwarding agents have deleted their copies first. This avoids mismatches of affinities across routers should a new affinity assignment request be received while a router still has the old affinity.

Service managers also keep affinities long enough after an outbound FIN is detected for a connection so that the final inbound ACK (or in the case of many Windows web browsers, the inbound RST) can be forwarded to the appropriate host. The use of a 'sticky' timer at the service manager satisfies this requirement. If a service manager changes an affinity at a time when it is possible that the affinity is still cached by a forwarding agent, the service manager asks the forwarding agents to delete the affinity before sending the updated affinity.

It should be noted that fixed affinities and wildcard affinities do not themselves include actions in the data structures described above. For flexibility, actions are defined separately but are included with fixed affinities or wildcard affinities in an affinity update message. The associated actions are stored along with the fixed affinity or wildcard affinity on service managers and forwarding agents. Whenever a fixed affinity or a wildcard affinity is referred to as being stored on a forwarding agent or a service

manager, it should be understood that associated actions may be stored with the affinity, whether or not such actions are explicitly mentioned.

Likewise, other items may be included in a stored affinity data structure. For example, the affinity may include a time to live when it is sent by a service manager.

- 5 When the affinity is received by a forwarding agent, the forwarding agent may compute an expiration time from the time to live and store the expiration time along with the fixed affinity.

An architecture that includes service managers and forwarding agents for providing network services has been described. A message protocol for sending  
10 messages from service managers to forwarding agents and for reporting activity and forwarding packets from forwarding agents to service managers has been disclosed as well.

Next, specific processes implemented on forwarding agents will be described. The forwarding agents allow the functionality provided by the service manager to be  
15 integrated into the routing infrastructure. Forwarding agents may be implemented on routers, switches, or any other network device that can intercept packets, analyze the packets, and perform whatever actions are required by a service manager to provide the relevant network service.

Figure 11 is a flowchart illustrating a process implemented on a forwarding agent  
20 for handling IP packets. The process begins at 1100 when a forwarding agent receives an IP packet from a network interface. In a step 1102, the forwarding agent searches for a

fixed affinity that matches the packet. If a fixed affinity is found, instructions have already been received from a service manager for handling packets in the flow that includes the current packet. Control is then transferred to a step 1104 where the forwarding agent processes the packet according to the fixed affinity that was found.

5           If the fixed affinity does not provide that the packet is to be forwarded or dispatched to a certain location, further processing may be done by the forwarding agent on the packet using other matching fixed affinities with different service priorities. The process of searching for a fixed affinity or a wildcard affinity may be repeated for different service priorities. Thus, different service managers having different service  
10       priorities can send wildcard and fixed affinities to forwarding agents and the forwarding agents process the packets for each service priority. The process ends at 1106.

          If no fixed affinity is found in step 1102, control is transferred to a step 1110 and the forwarding agent searches its wildcard cache to determine whether a wildcard has been received that matches the packet. If a wildcard match is not found, control is  
15       transferred to a step 1112 and the IP packet is forwarded. In certain systems, the default may be to drop packets that do not match a wildcard affinity or a fixed affinity instead of forwarding such packets. If a wildcard match is found in step 1110, then control is transferred to a step 1114 and the packet is processed according to any instructions provided along with the wildcard affinity.

20           In this manner, the wildcard affinity may be used to implement rule based redirection policies in the forwarding agent. When an action is specified or the dispatch

flag is set in the wildcard affinity, a forwarding agent will redirect the packet according to the general instruction or rule that is so contained in the wildcard affinity. In addition, the forwarding agent may gather statistics on the packet or report to the service manager if an appropriate instruction is contained or flag is set in the wildcard affinity. Thus, the  
5 wildcard affinity allows the service manager to distribute rule based criteria and actions to the forwarding agents that do not require the forwarding agent to forward the packets back to the service manager for more specific instructions. The overhead associated with doing so is avoided.

The service manager may change the criteria and actions specified in the wildcard  
10 affinities by simply sending new wildcard affinities to the forwarding agents. In addition, the statistics kept by the forwarding agents may be periodically sent to the service manager according to instructions contained in the wildcard affinity or otherwise retrieved by the service manager. The service manager may determine that the criteria and actions should be changed based on the information obtained from the forwarding  
15 agents.

For example, the service manager may send wildcard affinities that cause forwarding agents to distribute connections among a group of servers according to the IP address of a client making the connection. The forwarding agents may also be directed to keep statistics on the packets forwarded. The service manager can periodically obtain the  
20 statistics either directly or through an intermediary and then adjust the distribution of connections as required by sending new wildcard affinities to the forwarding agents. The forwarding agents and the service manager do not need to communicate or forward

packets back and forth for each connection. The forwarding agent handles a plurality of connections based on a rule or set of rules embodied in one or more wildcard affinities.

Figure 12 is a flowchart illustrating a process implemented on a forwarding agent to determine whether a packet matches a wildcard affinity. The process starts at 1200. In  
5 a step 1202, the forwarding agent determines whether the packet has the same protocol as the wildcard affinity. If the protocol is not the same, then control is transferred to a step 1212 where it is noted that there is no match and the process ends. If the protocol is the same, then control is transferred to a step 1204 where it is determined whether the affinity source IP address ANDed with the affinity mask is equal to the packet source IP  
10 address ANDed with the affinity mask.

If the two results are not equal, then control is transferred to step 1212. If the two results are equal, then control is transferred to a step 1206 where it is determined whether the affinity destination IP address ANDed with the affinity mask is equal to the packet destination IP address ANDed with the affinity mask. If the two results are not equal,  
15 then control is transferred to step 1212. If the two results are equal, then control is transferred to a step 1208 where it is determined whether the affinity destination port equals the packet destination port or the affinity destination port is zero. If neither is true, then control is transferred to step 1212. If either is true, then control is transferred to a step 1210 where it is determined whether the affinity source port equals the packet  
20 source port or the affinity source port is zero. If neither is true, then control is transferred to step 1212. If either is true, then control is transferred to a step 1211 where a match is noted and then the process ends.

In addition to receiving packets, finding affinity matches, and processing packets according to instructions included in the matching wildcard affinities or fixed affinities, the forwarding agent must also be able to process the wildcard affinities and fixed affinities received from service managers and store those affinities along with any associated actions so that the affinities and actions may be found when packets are received.

In addition to receiving IP packets and processing IP packets according to whatever wildcard and fixed affinities are stored on the forwarding agent, the forwarding agent must also receive affinities from a service manager. Received affinities must be managed, stored, updated and deleted when they expire. Figure 13 is a flow chart illustrating a process implemented on a forwarding agent for handling a wildcard affinity received from a service manager.

The process starts at 1300 when a wildcard is received. In step 1302, the forwarding agent looks up the wildcard in its database of stored wildcards to determine whether it matches an already stored wildcard. The forwarding agent may use a number of different data structures for storing wildcard and fixed affinities. In one embodiment, wildcard affinities are stored in a linked list and fixed affinities are stored in an AVL tree. It should be noted that in other implementations affinities may be stored in other data structures dictated by the particular needs of a given system.

If the wildcard affinity is found, then control is transferred to a step 1304 and the old wildcard is deleted. If the wildcard is not found, then control is transferred to a step

1306 and the forwarding agent checks the required actions listed in the wildcard affinity. In a step 1308, the forwarding agent determines if all of the listed actions are supported. If the listed actions are not all supported, then control is transferred to a step 1310 and the forwarding agent sends an affinity deny message to the service manager that sent the  
5 wildcard affinity. The affinity deny message identifies the actions included in the wildcard affinity that the forwarding agent does not support and gives the service manager the opportunity to attempt to send a different wildcard affinity with modified required actions. The process ends at 1312.

If the forwarding agent does support all of the required actions in the wildcard  
10 affinity, then control is transferred to a step 1314 and the forwarding agent stores the affinity in its wildcard data structure. The process ends at 316.

Figure 14 is a flow chart illustrating the process implemented on a forwarding agent for checking fixed affinities or wildcard affinities stored in an affinity data structure for the purpose of removing expired affinities. The process starts at 1400. The  
15 forwarding agent goes to the first affinity to be checked in a step 1402. The expiration time stored with the affinity is checked in step 1404. If the affinity has expired, that is, if the expiration time stored is before the current time, then control is transferred to a step 1406 and the affinity is deleted. If the expiration time is after the current time, then control is transferred to a step 1408 and the process checks whether the affinity is the last  
20 affinity to be checked.

If it is not the last affinity to be checked, then control is transferred to a step 1409 and the next affinity is accessed. Control is transferred back to step 1404 and the expiration time for that affinity is checked. If step 1408 determines that the last affinity has been checked, then the process ends at 1410. It should be noted that the process may  
5 simply repeat and go back to checking the first affinity or the process may end and be executed again by the forwarding agent after an appropriate interval.

The forwarding agent gathers statistics by intercepting packets, processing the packets to determine what statistics to record for the packet, and recording and reporting the statistics. Forwarding agents may record different types of statistics. In one  
10 embodiment, forwarding agents record the number of bytes of packets received corresponding to a flows that a service manager has designated for statistics gathering. As described above flows or sets of flows are designated for statistics gathering by setting the information flag in either a fixed or wildcard affinity.

In one embodiment, forwarding agents record the number of packets sent for each  
15 flow. Also, the forwarding agents may select whether to record number of bytes, number of packets, or some other detected quantity based on directions from a service manager. In addition to gathering statistics, forwarding agents also report the statistics gathered to a statistics collector. In one embodiment, the statistics collector is a service manager. In other embodiments, the statistics collector may be another device or else the statistics  
20 collector may be a special forwarding agent. Forwarding agents have a connection to the statistics collector for the purpose of sending reports. That connection may be the same network connection in which the forwarding agents receive packets or the same

connection which forwarding agents use to send packets to service managers. The connection may also be an independent connection.

A forwarding agent analyzes packets that match one of its stored fixed or wildcard affinities to determine how to classify the packet for the purpose of counting.

5 Figure 15 is a flowchart illustrating a process running on a forwarding agent for recording statistics about a packet. The process starts at 1500. In a step 1502, the forwarding agent receives the packet. The packet may either be received on the forwarding agent's network interface or the packet may be received from a service manager on the forwarding agent's service manager interface. The service manager may include a fixed  
10 affinity along with the packet or the forwarding agent may find a fixed or wildcard affinity that corresponds to a packet received on its network interface. In either case, an affinity is identified that matches the packet in a step 1503.

The affinity includes an information flag and a dispatch flag. The affinity also includes or may include a forwarding address field and a Network Address Translation  
15 (NAT) address field. The information flag is used by the service manager to specify whether statistics are to be gathered for the packet. In some embodiments, the information flag is one bit and in other embodiments, the information flag is more than one bit and may specify the type of statistics to be gathered. For example, the number of packets may be counted, or the total number of bytes in packets corresponding to a flow  
20 may be counted.

The forwarding agent checks whether the information flag in the affinity is set in a step 1504. If the information flag is not set, then control is transferred to a step 1506 and the forwarding agent forwards the packet without counting it. The process then ends at 1508. If the information bit is set, then control is transferred to a step 1510 and the forwarding agent checks whether the dispatch flag is set. The dispatch flag indicates whether the packet is to be dispatched. The packet is dispatched by the forwarding agent to the forwarding address specified by the affinity if the dispatch flag is set in the affinity. Packets are dispatched to the forwarding address regardless of what the destination IP address in the packet header is. It should be noted that, in certain embodiments, when the fixed affinity specifies that a packet destination IP address is to be translated to a NAT address, then the dispatch flag may be interpreted in another manner as is described below.

In one embodiment, the forwarding agent is distributing packets among a group of servers that service a virtual IP address. If the dispatch flag is not set and NAT is not being implemented, then the packet corresponds to an outbound packet because outbound packets do not need to be dispatched to an IP address other than the IP address specified in the destination IP address of the packet header. The servers include the correct IP address for the intended recipient when the packet is sent. The outbound packet is counted as an outbound packet in a step 1512 if it is determined that the dispatch bit is not set in step 1510. After the packet is counted as an outbound packet in step 1512, control is transferred to a step 1514 and the packet is forwarded. The process then ends at 1516.

If, in step 1510, the dispatch flag is set, then control is transferred to a step 1520 and the packet is counted as an inbound packet. Inbound packets are dispatched to the forwarding address provided in the fixed affinity so that packets bound for a virtual IP address may be directed to a real machine. Once the inbound packet is counted, control  
5 is transferred to step 1514 and the packet is forwarded.

If NAT is specified, then the fixed affinity will include a NAT address to be used for translating the destination IP address of certain packets. In such a case, the information flag is still used to indicate that packets are to be counted and the dispatch flag is used in one embodiment to indicate whether the packet is an inbound packet or an  
10 outbound packet. In some embodiments, the source and destination IP addresses may be used to determine whether the packet is inbound to a group of servers or outbound from a group of servers. However, using the dispatch flag for this purpose may speed up that determination. In one embodiment, the dispatch flag being set indicates that the packet is inbound and the dispatch flag not being set indicates that the packet is outbound. Packets  
15 are handled by the forwarding agent in a similar manner to that described in connection with Figure 15, except that the destination IP addresses and possibly the source IP addresses are translated and the packets are forwarded normally and not dispatched.

Figure 16A is a table illustrating the information and dispatch flags and the forwarding address field for a fixed affinity. If the forwarding address field contains a  
20 forwarding address and the dispatch flag is set but the information flag is not set, then the packet is dispatched without being counted. If the information flag is set, then the packet is counted as an inbound packet. If the information flag is set, but the dispatch flag is not

set, then the packet is counted as an outbound packet. In one embodiment, if the forwarding address is null, but the dispatch flag is set, then the packet is dropped. The information flag may be set or not set indicating whether or not the packet is to be counted before it is dropped. In different embodiments, the forwarding address may be set to a special null address or may simply not be written to indicate that the packet is to be dropped.

Figure 16B is a table illustrating different values for the NAT address field and the information and dispatch flags for fixed affinities that specify network address translation as an action. If network address translation is specified, then a NAT address is provided in the fixed affinity. If the information flag is set, then the packets which are translated are counted. The dispatch flag is used to indicate whether the packet is inbound or outbound. It is assumed that if NAT is specified, that the packet will not actually be dispatched to an address different than the one specified in the packet header after address translation.

A generalized method of sending instructions from a service manager to a forwarding agent so that the forwarding agent can gather statistics about packets being serviced has been disclosed. The forwarding agent does not need to be configured with a statistics gathering scheme for any specific application. The forwarding agent processes packets for the purpose of gathering statistics according to instructions received by the service manager. An efficient protocol for providing such instructions and interpreting them has been described.

It should be noted that in different embodiments, different methods of storing and accessing expiration times of affinities may be used. In one embodiment, the time to live included in a fixed affinity or wildcard affinity that is received by the forwarding agent is added to the current time when an affinity is received. The result is an expiration time  
5 that is then stored in the forwarding agent's affinity data structure for the purpose of providing an expiration time used by an expired affinity deletion process such as is described in Figure 16. In other embodiments, the time to live along with a current time stamp may be stored in the affinity data structure. Other arrangements for tracking affinity expiration times such as using a table that includes pointers to affinities along  
10 with expiration times may be used. Deletion of affinities may be optimized according to the requirements of a specific system.

Processes have been described that illustrate how forwarding agents handle messages received from service managers and manage their affinity databases. Also, it has been shown how forwarding agents check for wildcard affinity matches when an IP  
15 packet is received over a network interface. Thus, forwarding agents are able to receive general wildcard affinity instructions from service managers and process packets according to those instructions to provide distributed, rule based network services.

Up to this point, an embodiment has been described wherein a single tier of forwarding agents receives instructions from a service manager and forwards packets to a  
20 single tier of servers. In some embodiments, it is desirable to forward packets and balance workload in a serial fashion among two or more tiers of servers. In one embodiment, each tier of server provides a different service. In the example described

below, requests from a client are first filtered through a tier of firewalls before access is allowed to a tier of web servers. Significantly, the same firewall is used for both inbound and outbound flows.

Figure 17 is a block diagram illustrating a multi-tiered processing scheme. A client 1702 is connected via a network 1704 such as the Internet to a group of network appliances shown as web servers 1706. In between client 1702 and the web servers are multiple tiers of forwarding agents and an additional tier of network appliances shown as firewalls 1708. A first tier of forwarding agents 1710 is interposed between the client and the firewalls 1708. A second tier of forwarding agents 1712 is interposed between the firewalls and the web servers 1706. It should be noted that although the forwarding agents are shown separately from the network 1704, since the forwarding agents are preferably implemented on routers, one might also consider the forwarding agents to be integrated into or part of the network.

The first tier of forwarding agents 1710 is distinct from the second tier of forwarding agents 1712. Forwarding agents 1710 do not forward packets directly to web servers 1706. Likewise, the forwarding agents in the second tier are connected to web server 1706 but they only receive packets from firewalls 1708 and do not receive packets directly from the client that are not forwarded by the first tier of forwarding agents 1710. Both tiers of forwarding agents, however, are in communication with a common service manager (not shown) for reasons described below. In one embodiment, the first tier or forwarding agents is recognizable and distinguishable from other tiers of forwarding agents by some convention such as the control addresses of the forwarding agents being

in the same subnet or the ports being the same. The forwarding agents may also be distinguished by configuring them to the service manager. In some embodiments, there may be more than one service manager for a tier. In that case, the service managers may share data with each other so that each service manager has access to the list of

5 forwarding agents in a tier.

When a client sends a request to one of the web servers, the request is first intercepted by one of the forwarding agents in the first tier. That forwarding agent forwards the request to the service manager. The service manager assigns the request to one of the firewalls and that firewall, after performing whatever authentication or checks  
10 are appropriate, then forwards the request to one of the forwarding agents in the second tier. The forwarding agent in the second tier then forwards the request back to the service manager.

The service manager identifies the forwarding agent as being a forwarding agent in the second tier and sends instruction to the forwarding agent to forward the packet to  
15 one of the web servers. Thus, the service manager sends different instructions to different forwarding agents based on the tier of the forwarding agent. In one embodiment, the service manager builds a connection object that notes the IP address of the firewall chosen for an incoming packet received from the first tier of forwarding agents. That connection object is used to make sure that an outgoing packet forwarded  
20 by the second tier of forwarding agents is forwarded to the same firewall.

Thus, the request from the client and the server response to the request are forwarded through the same firewall. The service manager receives the client request from the first tier of forwarding agents and routes that request to an appropriate firewall. Since that assignment is stored in a connection object, the service manager can insure that  
5 the same firewall is assigned to the response to the request sent by server through the second tier of forwarding agents. The use of a common service manager enables such state information about a connection to be stored and allows the forwarding decision implemented by the second tier of forwarding agents to be dependent on the forwarding decision that was made for the first tier of forwarding agents.

10 It should be noted that the designation of client and server in Figure 17 is arbitrary and is provided for the purpose of illustration only. Likewise, the designation of the first tier of network appliances as firewalls and the second tier network of network appliances as web servers is also made for the purpose of example. In general, a service manager and multi-tier forwarding agent architecture may be used with any network appliances  
15 and with any relationship between the two end parties that are communicating.

As described above, the service manager keeps track of which forwarding agents are included in each tier of forwarding agents and also keeps track of which network appliances are included in each tier of network appliances. In addition, the service manager keeps track of which servers in the various tiers have been assigned to various  
20 connections.

Figure 18A is a block diagram illustrating a data structure in the service manager that is used to keep track of forwarding agents in the first tier. Consistent with the example shown in Figure 17, the tier one forwarding agents include forwarding agent 1 and forwarding agent 2.

5        Figure 18B is a block diagram illustrating a data structure used to keep track of the forwarding agents in a second tier. The second tier forwarding agents include forwarding agent 3 and forwarding agent 4. In addition, other tiers of forwarding agents may be included as corresponding data structures to keep track of each tier of forwarding agents.

10       Figure 18C is a block diagram illustrating a data structure to keep track of the appliances in the first tier. The first tier appliances include firewall 1, firewall 2, and firewall 3.

Figure 18D is a block diagram illustrating a data structure for keeping track of the appliances in the second tier. The tier two appliances include web server 1, web server 2,  
15       and web server 3. In addition, other data structures may be included to keep track of additional tiers of appliances. It should be noted that the data structures may be implemented in any appropriate way including as individual tables, as one combined table or as various forms of linked lists.

Figure 18E is a block diagram illustrating a connection object data structure for  
20       keeping track of servers assigned to various connections. The connection objects include a protocol, a client IP address, a client port, a server IP address, and a server port. In

addition, the connection object includes a tier one server along with additional tier servers depending on the number of tiers included in the system. It should be noted that the designation client and server again are arbitrary and the client IP address and port are merely meant to specify the IP address and port of the packets that are initially received  
5 by the first tier of forwarding agents when a connection is being requested. Depending on the direction of the flow, the client and server addresses and ports are designated as source or destination in the packet.

Thus, the service manager stores the information regarding the forwarding agents that are in each tier and the appliances that are in each tier. In addition, the service  
10 manager also notes for each connection the server in each tier that is selected to handle that connection. When a packet is forwarded to the service manager by the forwarding agent, the service manager determines what tier the forwarding agent belongs to and selects an appropriate network appliance in that tier to handle the packet. In addition, if the packet corresponds to an outgoing flow, the service manager makes sure that the  
15 same network appliance that handled the incoming flow is assigned to the packet belonging to the outgoing flow. Again, it should be noted that the designations of incoming and outgoing are purely arbitrary and are intended only for the purpose of example.

Figure 19 is a flowchart illustrating a process implemented on the service  
20 manager when a packet is received from a forwarding agent. The process starts at 1900 when the packet is received. In step 1902, the tier of the forwarding agent is determined. If the forwarding agent is in the first tier, control is transferred to step 1904 and a first

tier appliance is selected to handle the connection. The selected first tier appliance is stored in a connection object. Control is then transferred to step 1906 and the instructions are sent to the first tier forwarding agents. In some embodiments, instructions may only be sent to one forwarding agent instead of the entire first tier. In step 1908, the packet is  
5 forwarded to the selected forwarding agent. It should be noted that in some embodiments, steps 1906 and 1908 may be combined. The process then ends at 1916.

If, in step 1902, it is determined that the forwarding agent belongs in the second tier, then control is transferred to step 1910 and a second tier network appliance is selected to handle that connection. Control is then transferred to step 1912 and  
10 instructions are sent to the second tier forwarding agents for handling packets included in the flow of received packets. The packet is then forwarded in step 1914 and the process ends at step 1916.

Figure 20 is a flowchart illustrating a process implemented on the service manager for insuring that outgoing packets are assigned to the same firewall as incoming  
15 packets belonging to the same connection. This is important for various types of network appliances like firewalls that keep track of connection state information. It is important that all packets in both flows which comprise the connection be routed through the same network appliance. The process starts at 2000. In step 2002, it is determined whether the forwarding agent tier directs packets to a firewall. If the forwarding agent does not direct  
20 packets to a firewall, then control is transferred to step 2004 and the service manager selects a network appliance from among the network appliances included in the tier that is connected to the forwarding agent. In one embodiment, a load balancing algorithm is

used to select a network appliance. In another embodiment, a predetermined scheme is used to select a network appliance. The process ends at 2006.

If, in step 2002, it is determined that the forwarding agent tier of the forwarding agent is connected to a firewall or other network appliance that keeps track of a connection state, control is transferred to step 2010 where it is determined whether a flow has been established for the forwarding agent tier that is on the other side of the firewall. If a flow has not already been established, then control is transferred to step 2004. If a flow has been established, then control is transferred to step 2012 and the intercepted packet flow is assigned to the same firewall as the flow established in the other direction. The process then ends at 2014.

Thus, when the service manager receives a packet from a forwarding agent, it first checks whether that forwarding agent is connected to a firewall. If the forwarding agent is connected to a firewall, then the service manager checks whether a flow has already been established from a forwarding agent on the other side of the firewall and then assigns the current packet flow to the same firewall. As noted, other types of network devices that maintain the state of a connection may receive the same treatment as firewalls.

Forwarding packets and balancing workload in a serial fashion among two or more tiers of servers has been disclosed. In the example described above in detail, requests from a client are first filtered through a tier of firewalls before access is allowed to the tier of web servers. A service manager keeps track of the network topology by

storing information about the tier of each forwarding agent. Packets forwarded from each tier of forwarding agents are appropriately directed to either a firewall or a web server. The service manager also ensures that the same firewall is used for both inbound and outbound flows.

5           Although the foregoing invention has been described in some detail for purposes of clarity of understanding, it will be apparent that certain changes and modifications may be practiced within the scope of the appended claims. It should be noted that there are many alternative ways of implementing both the process and apparatus of the present invention. Accordingly, the present embodiments are to be considered as illustrative and  
10       not restrictive, and the invention is not to be limited to the details given herein, but may be modified within the scope and equivalents of the appended claims.

WHAT IS CLAIMED IS: